EXPORING THE VALUE OF SENSORS TO A RECCE UNIT USING AGENT-BASED MODELS

DR MICHAEL K LAUREN
DTA, NZ

LT COL DEAN L BAIGENT
CAD BRANCH, NZ ARMY

Reprinted here, courtesy of the Journal of Battlefield Technology
(www.adfa.edu.au/jbt)

EXPLORING THE VALUE OF SENSORS TO A RECCE UNIT USING AGENT-BASED MODELS

M. K. Lauren and D. L. Baigent

Abstract. The increasing presence of electronic warfare devices on the battlefield, combined with doctrine that emphasises the value of manoeuvre, precision strike, and high operational tempo, present the military operations researcher with increasingly difficult problems. Additionally, it is becoming increasingly obvious that traditional OR techniques have serious limitations for describing complex systems with non-linear dependencies, which are common on the battlefield. This paper describes how a cellular automata model can make some headway on the problem of describing the modern reconnaissance environment. The model emphasises the behaviour of the participants rather than the physics of the equipment. This leads to complex interactions between the entities, which go some way towards representing the non-linearities inherent in real-life operations. The value of detection-range advantage and aerial reconnaissance falls out of the model remarkably naturally when one considers the arbitrary way these are represented in conventional models. Moreover, it is seen that for certain ranges of parameters, the survivability of the Recce force is nearly independent of the kill probability of the weapons of its opponents, a result that contrasts with the Lanchester-like nature of conventional models. For these reasons, the results presented should be of great significance to the military OR community.

INTRODUCTION

Background

The New Zealand Army has begun a process of operational research to determine appropriate force structures for 2005 and beyond. To support this process, the Army and the Defence Force science organisation, the Defence Technology Agency, have established a permanent operations analysis cell.

Though New Zealand is a relative latecomer to the establishment of such a capability (although various organisations within NZ have conducted operational research as required), the timing of NZ's involvement in the field has fortuitously coincided with a shift in emphasis in the OR community. Increasingly, there is a recognition that the range of conventional models used for such studies (including detailed physics-based combat models), are often inadequate for describing the conditions under which real forces must operate [1–3]. This is generally attributed to the "non-linear" nature of the world. An example of this is the so-called butterfly effect, where, anecdotally, the flapping of a butterfly's wings is supposed to have a profound effect on the weather on the other side of the world. Given the apparent critical importance of boundary conditions to non-linear problems, obtaining useful output from a model appears almost hopeless.

However, recent developments in the fields of physics, chemistry and biology have led to the evolution of models which at least offer some hope for this kind of analysis [4]. These advances are embodied in a rapidly evolving field called complexity theory, which might best be described as a melting pot of statistical physics, numerical and Monte Carlo methods, biological and genetic models and metaphors, and scaling and renormalisation group methods.

Initial attempts to incorporate these methods into our analysis [5] led DTA to become involved with the US Marine Corp senior science advisor's attempts to explore the same issues [6]. DTA's involvement in the USMC Combat Development Command's Project Albert has allowed it access to an attempt at modelling the non-linear nature of combat, called ISAAC (Irreducible Semi-Autonomous Adaptive Combat) [7]. This model concentrates on behaviour of the participants, rather than the detailed physics of their equipment. The complex interactions between the entities in the simulation tend to be quite non-linear [8], with the emphasis of the analysis shifting from questions of the theoretical capabilities of a piece of equipment, to the possible alternate ways in which war or peacetime operations might be conducted given certain capabilities.

Certainly, the increasingly sophisticated electronic equipment employed by modern land forces, has led to a dramatic change in the way conflicts are conducted and armies are structured. Although the value of such equipment has been demonstrated, most famously in the Gulf War, it is much harder to incorporate the value of a sensor as opposed to a gun system into a conventional analytical combat model. The reason for this is that the benefit is not embodied in the firepower such a piece of equipment adds to a unit. Rather, it is in how such an "informational" edge can be exploited. In terms of conventional combat models, such an edge fits in the category of an "intangible", where we use this term to mean something that cannot be described by physics alone.

Question at Hand

As part of its OR effort, the New Zealand Army considers a theatre of operation which sees a motorised battalion operating as part of an ABCA (America Britain Canada Australia) brigade in mid-intensity conflicts. The battalion should not, by choice, participate in the main engagement with heavy forces, since it lacks the firepower and protection of Main Battle Tanks or similar vehicles. The battalion has been assigned a motorised reconnaissance capability. One concept of operations for this reconnaissance force is based on the hypothesis that it should "steal" information (that is, stealthily monitor its opponents).

The principal conventional models available in New Zealand to test the feasibility of this approach on the modern battle are the US wargame model JANUS, and the UK-developed analytical model CAEN. However, attempts to test such a hypothesis with these models only served to highlight their limitations. To see why, consider the two modes in which such models are used. The principal method of analysis involves setting up the model with certain pre-programmed routes and orders, with limited ability to adapt to the circumstances of a particular run. Once set up, the model is run multiple times. It falls to the analyst to set up the scenario in such a way that it represents how the participants behave.

But how does the analyst go about this? Assume the Blue force has superior sensors to the Red. Should the scenario be set up so that Red never detects Blue (because Blue can see Red coming)? Or does the analyst allow Red to come close enough to Blue to detect it at certain times? Is Red too stupid to do anything to counter Blue's superiority? Is Blue's superiority sufficient that it can't be overcome?

Essentially, the problem is that the analyst has 100 per cent knowledge of each force, but can only incorporate this knowledge into each side's behaviour in an arbitrary way. Furthermore, the runs of the scenario only vary from each other by a limited degree. Such a model resembles a stochastic Lanchester equation, with a certain portion of each force becoming engaged, and the outcome depending just on the number of these engaged entities and their kill probabilities.

The second approach is to use the model as a "game" with human players. It is thus up to each player to exploit the sensor advantage. Using this approach it soon becomes clear that the players require extremely quick reactions to make sure all their vehicles react in a sensible way having detected the enemy. It is virtually impossible to monitor all the vehicles closely enough to stop them blundering past the point at which they acquire the enemy and into range of that enemy's sensors.

An alternative approach presented in this report is to use multiple "agents". This approach is slowly but surely finding broadening support within the military OR community [9,10,11]. An agent means an entity within the model which makes its own decision on how to behave based on its circumstances. The framework used here is known as a cellular automata model. Standard models of this kind use a grid of cells, which may at most contain a single automaton. The model uses generic units (the automata) which are described in terms of simple capabilities and personalities, rather than specific pieces of equipment.

Owing to ISAAC's simplicity, it is possible to model a wide range of behaviour quickly and in a systematic way. This allows the model to be used with the intention of exploring the feasible range of solutions to a given problem, rather than producing a single supposed "answer".

The ISAAC Model

ISAAC is relatively simple and for the sake of brevity, only a short description is given. It is described in detail in Ilachinski [7], and the executable program may be downloaded from the Web address www.cna.org/ISAAC.

The automata parameters break down into three classes: attributes, personalities and meta-personalities.

- The first class describes movement rate, weapons range, kill probability, sensor range and so on.
- The second class comprises weightings that describe an automaton's propensity to move towards/away from friendly/enemy automata, and towards/away from a goal (flag) point. The model calculates automata moves by summing the number of friendly and enemy automata within a threshold range of each square within movement range, and uses the personality weightings to determine the "penalty" associated with each move in terms of how it positions the automaton relative to the friendly/enemy forces and the goal point. Communications may also be simulated, by including the number of friendly/enemy automata visible to other friendly automata within a set communications range.
- The third class modifies the above procedure. Three parameters fit this category: the cluster parameter, which "turns off" an automaton's propensity to move towards friendly automata once a threshold cluster size has been reached; the advance parameter, which requires a threshold number of friendly automata to be surrounding an automaton before it will advance towards its goal; and a combat parameter, which will only allow an automaton to advance towards the enemy once a threshold numerical advantage has been achieved.

The automata themselves can be in one of three states: *alive*, in which case they use the baseline parameters; *injured*, where a secondary set of parameters may be used to indicate the automata has suffered damage; and *dead*, in which case the automaton is removed from the battlefield.

RECONNAISSANCE SCENARIO

To test the hypothesis in the introduction, a scenario was designed around the reconnaissance of a high-value target and the protection of the same target by a counter-reconnaissance force. The Blue force conducting the reconnaissance is based on four pairs of wheeled vehicles (four Blue "dots"). The target is defended by eight pairs of wheeled vehicles (eight yellow dots), with the same capabilities as Blue, but in the surveillance/early-warning role. A quick reaction force employing tracked armoured vehicles further protects the target. It also consists of eight pairs of vehicles (eight Red dots). The scenario uses an 80x80-cell grid without terrain features. Note that the automata only "notionally" represent these vehicles. They could as easily represent squads of infantry with similar capabilities and operational behavior.

To provide a framework within which to choose realistic values for the automata capabilities, it was imagined that the model represented an area of about 40km x 40km, and each time step about 1 minute. The reconnaissance force was assumed to have light weaponry, equivalent to a heavy machinegun, with an effective range of about 0.5km (hence 1 square). The Red reaction force was assumed to have a much more potent weapon, something similar to a 75mm gun, with an effective range of at least 1.5km (3 squares).

In addition to these capabilities, the automata are given "personalities". For this scenario, these were:

Blue:

- always retires from Red;
- does not require the support of other Blue to advance to goal;
- units have the ability to communicate the location of Red units to each other; and

- driven to reach the goal with different levels of aggression.

Red:

- will aggressively chase Blue;
- use communications to concentrate fire on Blue;
- injured Red have slight propensity to stay near other Reds;
- will attack Blue without requiring support; and
- yellow units do not chase Blue, but will engage Blue if not outnumbered.

Figure 1 shows a sample run. The parameters used for the run are given in Table 1. Note how a pair of Blue units have drawn in the defenders, creating a gap for the others to exploit. The final outcome is a single Blue unit lost, while another Blue unit "discovers" the goal.

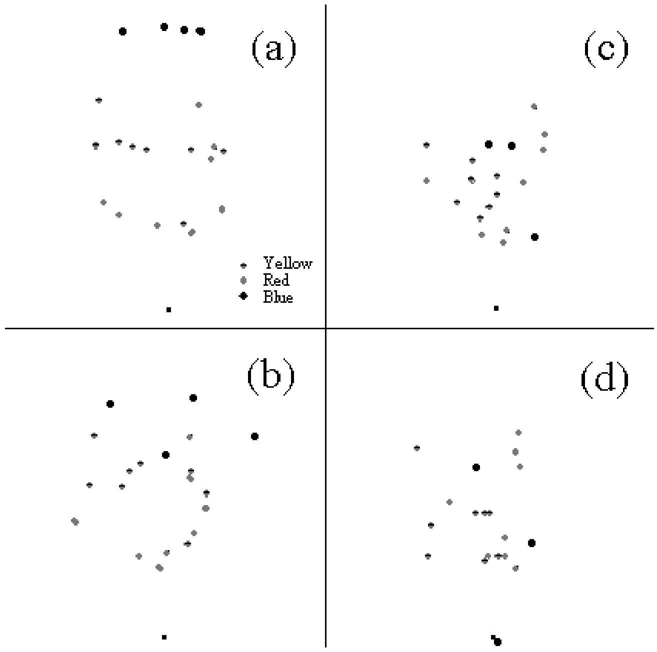

Figure 1. A single run. The Blue force starts at the top of the grid. The black square at the bottom is the location of the high-value target. Snapshots (a) through to (d) show progression of the run.

RESULTS

Fitness Landscapes

Results were produced in the form of a "fitness landscape". This is a three-dimensional plot which indicates the performance of either the Blue or Red force for various combinations of parameters. Each point on the landscape is the average of 100 runs of the scenario. The higher the landscape, the better the performance.

The measure of effectiveness used in this paper is number of surviving forces. However, consideration was given to time taken to reach goal, so that the results were meaningful (clearly, if Blue sits in the top right-hand corner, it is less likely to suffer casualties, but is not achieving its objective). Blue units have a weighting of 40 (note that weightings are out of 100) for retreating from Red. Thus, as the goal weighting increases above 40, the Blue force will increasingly become inclined to move towards the goal rather than retreat from Red. It was found that increasing aggressiveness did not have a dramatic effect on the average time taken to reach the goal. For example, increasing the goal weighting from 10 to 90 for cases where Blue had a detection range advantage decreased average time to goal by only 20%, but increased casualties by more than 30%.

Importance of Detection Range Advantage

Of particular interest is the value of detection capabilities compared with firepower. "Sensor range" here is treated as the range at which a given automaton can detect an enemy, for whatever reason. Thus it represents a combination of factors, such as stealth, training/experience, and sensor capabilities. Figure 2 illustrates the effects on Blue survivability of varying both the kill probability of Blue's weapons and the range at which Blue can detect Red. Additionally, for this particular data plot, Blue fire range was extended to 1.5km, since increasing Blue firepower would otherwise have limited effect if Red could outrange Blue.

It can be seen that once a detection range advantage is gained (Red has a detection range of 3), there is a dramatic jump in Blue survivability. By contrast, there is only a gradual improvement as firepower increases. It is also notable that increasing firepower does not make much difference to Blue survivability until Blue has achieved a significant firepower

Threshold: 2	Personalities: weighting towards			Meta Personalities			Attributes			
	Red	Blue	Goal	Cluster	Advance	Combat	Ranges: fire	Sensor	Movement	Kill prob.
Blue (alive/injured)	-40/-40	-10/0	variable	2/2	0/0	1/10	1/1	variable	2	0.1
Red	10/10	5/50	0	0/0	0/0	-4/-2	3/3	3	2	0.3
Yellow	5/5	0/0	0	0/0	0/0	0/0	1/1	3	2	0.1
Aerial	-5/-5	10/10	0	0/0	0/0	-10/-10	0/0	25	4	0.0

Table 1. Model parameters.

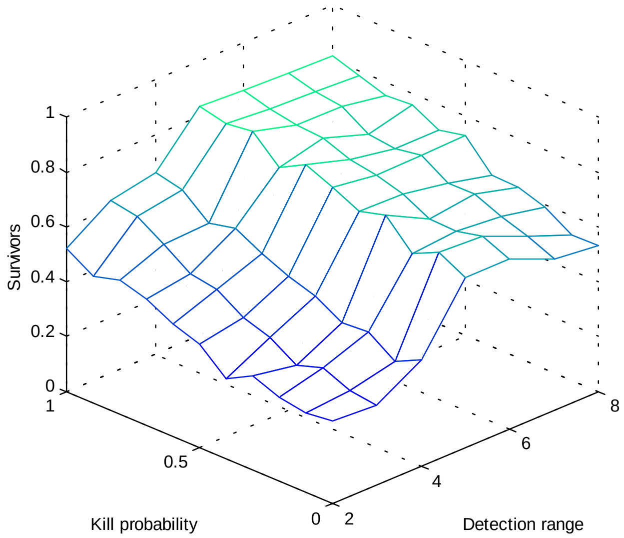

Figure 2. Effect of varying Blue kill probability and detection range on number of Blue survivors.

advantage (Red kill probability is 0.3).

Also significant is the effect of improving Red firepower and numbers, illustrated in Figure 3 for a case where Blue has a detection range advantage. Here, increasing Red firepower has almost no effect on Blue survivors above a kill probability of about 0.2.

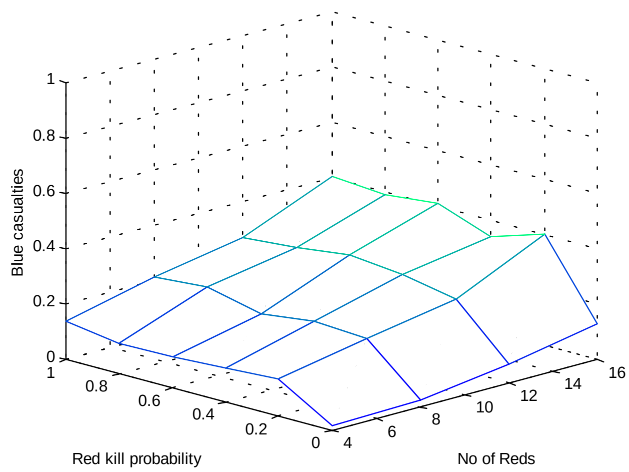

Figure 3. Effect of increasing Red kill probability, and increasing number of Red units, for a case where Blue has a detection-range advantage.

Note that this is a comment on Blue protection requirements, implying that Blue protection levels make no difference to Blue survivability, if Red has a kill probability of at least 0.2. The reason why protection levels make little different once Red achieves this level of firepower is that Red is most hampered by having to locate Blue units. Once found, Red can use concentrated fire to dispatch Blue units relatively easily, even if individually Red has moderate firepower [12].

This is in stark contrast to a Lanchester model, for which kill probability should have a strong effect on Blue casualties.

Aerial Observer

Red's ability to communicate Blue positions to other nearby Red units plays an important role in its ability to concentrate firepower. This suggests that an aerial observer (with superior detection range to ground-based units) relaying this information should produce a significant improvement in Red's ability to destroy Blue.

To explore this possibility, extra automata were added to represent aerial observers of some kind. The aerial observer was given a detection range of 25 cells, could move four cells at a time, and could communicate the position of Blue units to Red units within 25 cells. Note the detection range assumed represents a vehicle with capabilities that are towards the top-end of aerial sensing capabilities.

Figure 4 shows the effect of the presence of Red aerial observers on Blue survivability, for a case where Blue has a detection range advantage over Red ground units. Clearly, as soon as the aerial observer can communicate with other Red units, Blue casualties climb dramatically. Increasing the number of aerial observers improves performance further, but the improvement diminishes as the number added increases, since it only takes the presence of a small number of such observers to cover the entire battlefield.

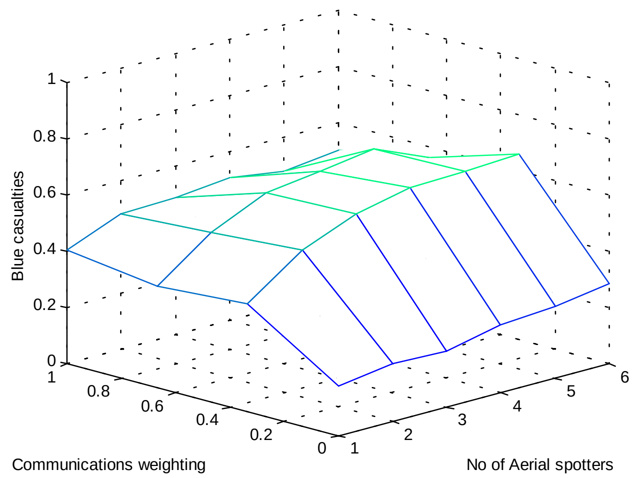

Figure 4. Influence of Red aerial observer on Blue casualty levels. Communications weighting of 0 means the aerial observer has no influence.

Note the seemingly contradictory drop in Blue casualties for the highest values of Red communications weighting. At this point, the Red units have so much weighting placed on the information supplied by the aerial observers that they are immediately drawn to the observer whenever a Blue unit is spotted, leaving gaps for the remaining Blue units. That is, Red behaves stupidly for these values.

SUMMARY AND CONCLUSIONS

The ISAAC model provides many interesting insights into the value of detection-range enhancers. Specifically, the results point to significant improvements in the performance of either side if detection range advantages can be gained. The improvement is not only significant, but it is non-linear. In this case, the non-linearity is represented by a jump in Blue survivability if a detection range advantage exists.

The necessary threshold advantage is likely to be related to the speed at which the forces move. For example, if a Red unit is moving sufficiently quickly that it moves to a range at which it can detect Blue before Blue has had enough time to react to the detection of Red, then any detection range advantage that Blue has will be negated.

Whereas a significant jump in survivability can be achieved once a detection-range advantage is gained, increasing the armour or weaponry of the reconnaissance force produces a much more gradual improvement in survivability, meaning that significant advantages are required to improve survivability markedly. It is likely to be impractical to improve a reconnaissance vehicle's armoury to the necessary degree if it were facing a threat equivalent to a modern infantry fighting vehicle.

Bearing these things in mind, a feasible solution might be to use a low-signature vehicle operating with a great deal of stealth with an experienced crew, and which has both good sensor capabilities, and sufficient mobility to be able to get out of trouble quickly.

Alternately, a larger reconnaissance force may be able to combat the counter-reconnaissance force more effectively, since the reason why increasing firepower has little effect for the reconnaissance force is that the counter-reconnaissance force outnumbers it substantially. In particular, since Red is able to co-ordinate its firepower, it will destroy any encountered units quickly even when its weapons have modest performance.

The modelling of the presence of a Red aerial observer had a significant impact on the survivability of the reconnaissance force. A reconnaissance force without the capability to hide from an aerial observer is likely to suffer significant casualties against a counter-reconnaissance force of the strength modelled.

The results presented here are certainly not the definitive "answer" to the question of how reconnaissance should be conducted. Thought needs to be given to which other aspects and variations should be modelled. We saw that the inclusion of an aerial observer significantly altered the results. Are there other elements that should have been modelled which have a similar effect? And what would happen if Red behaved in a different manner?

For example, Red might opt to spread its surveillance units out in a uniform, static line across the bottom of the battlefield, so that their fields of view overlap. This way, Blue could not cross without being detected. Of course, if Blue detected such a line first, it could either circumvent it by going around it (by going outside the battlefield area), or by calling in strike to hit one of the static targets to make a hole in the line. Such counters and counter-counters to each side's tactics must be considered within the realms of what is feasible. ISAAC appears to be a useful tool for exploring these ideas.

REFERENCES

[1] J. Rhodes (Lt Gen), in *Maneuver Warfare Science 1998*, (ed) F. Hoffman and G. Horne, US Marine Corp Combat Development Command, 1998

[2] R. Scales (Maj Gen), "Adaptive Enemies: Achieving Victory by Avoiding Defeat", *Joint Forces Quarterly*, Vol. 23, Fall, 1999.

[3] J. Watson (Lt Col), "Living on the Edge", *The British Army Review*, Vol. 123, pp. 19-29, 1999.

[4] M. Waldrop, *Complexity: The Emerging Science at the Edge of Order and Chaos*, Simon and Schuster, New York, 1992.

[5] M. Lauren, "Modelling Combat Using Fractals and the Statistics of Scaling Systems", *Military Operations Research, Warfare Analysis and Complexity Special Issue*, in production.

[6] A. Brandstein, "Operational Synthesis: Applying Science to Military Science", *Phalanx (The Bulletin of the Military Operations Research)*, Vol. 32, No. 4, pp. 1, 30-31.

[7] A. Ilachinski, *Irreducible Semi-Autonomous Adaptive Combat (ISAAC): An Artificial-Life Approach to Land Warfare (U)*, Centre for Naval Analyses, CRM 97-61.10, 1997.

[8] M. Lauren, *Characterising the Difference Between Complex Adaptive and Conventional Models*, DOTSE Report 169, NR 1345, 1999.

[9] F. Hoffman and G. Horne (eds), *Maneuver Warfare Science 1998*, US Marine Corp Combat Development Command, 1998.

[10] C. Hunt, "Uncertainty Factor Drives New Approach to Building Simulations", *Signal*, Vol. 53, No. 2, pp. 75-77, 1998.

[11] A. Woodcock, L. Cobb, and J. Dockery, "Cellular Automata: A New Method for Battlefield Simulation", *Signal*, pp. 41-50, Jan 1988.

[12] M. Lauren, *Firepower Concentration in Cellular Automata Models—An Alternative to the Lanchester Approach*, DOTSE Report 172, NR 1350, 2000.

Dr Michael Lauren is head of operational analysis with DTA, NZ. His research interests are in the scaling properties of high-dimensional chaotic systems and complexity theory generally. Lt. Col. Dean Baigent commands the NZ Army's Capability, Analysis and Doctrine (CAD) Branch. CAD and DOTSE are jointly conducting OA to determine the Army's future requirements. Contact: michaell@dotse.mil.nz.